

What is claimed:

1. A computer-implemented method of searching a plurality of queries against at least one database containing a plurality of records, comprising the steps of:

- a. partitioning the plurality of queries into a set of smaller subsets of
5 queries;
- b. partitioning the at least one database into a set of smaller
subdatabases;
- c. designating searching tasks to be performed by associating each of
said subsets of queries with one or more of said subdatabases, assigning each searching task to
10 one of a group of computers operating in parallel, wherein each member of the group of
computers operating in parallel has at least one searching task assigned thereto, and executing at
least some of the assigned searching tasks using the group of computers operating in parallel;
and
- d. collecting search results from the executed searching tasks and
15 generating a unified search result in accordance with the collected search results;
wherein the partitioning of the queries and the partitioning of the database
are done by one or more members of the group of computers operating in parallel; and
wherein step c further comprises dividing at least one of the searching
tasks into two or more smaller searching tasks, and designating the two or more smaller tasks as
20 related tasks on a virtual shared memory bulletin board.

2. The computer-implemented method of claim 1, wherein the partitioning of the

queries and the partitioning of the database are done by each member of the group of computers operating in parallel.

3. The computer-implemented method of claim 1, wherein the partitioning of the queries and the partitioning of the database are based on the processing capacity of each member
5 of the group of computers operating in parallel.

4. The computer-implemented method of claim 1, wherein each member of the group of computers operating in parallel assigns to itself which searching tasks it will perform.

5. The computer-implemented method of claim 1, wherein the plurality of queries are searched against two or more databases.

10 6. The computer-implemented method of claim 5, wherein at least one database is not split into a set of smaller databases.

7. The computer-implemented method of claim 1, wherein at least one of the group of computers operating in parallel performs two or more searching tasks during executing of the searching tasks in step (c).

15 8. The computer-implemented method of claim 7, wherein when one of the members of the group of computers operating in parallel finishes a searching task, it assigns to itself another searching task during executing of the searching tasks in step (c).

9. The computer-implemented method of claim 1, wherein all of the assigned searching tasks are performed by the group of computers operating in parallel.

20 10. The computer-implemented method of claim 1, wherein each member of the group of computers operating in parallel is identical.

11. The computer-implemented method of claim 1, wherein at least two members of the group of computers operating in parallel are different.

12. The computer-implemented method of claim 1, wherein each member of the group of computers operating in parallel has the same operating system.

5 13. The computer-implemented method of claim 1, wherein at least two members of the group of computers operating in parallel have different operating systems.

14. The computer-implemented method of claim 1, wherein failure of one or more of the members of the group of computers operating in parallel does not affect the correctness of the search results.

10 15. The computer-implemented method of claim 1, wherein during execution of each assigned task in step (c), a numerical raw score is reported as part of the search result associated with such task, wherein the numerical raw score corresponds to a quantitative measure of a match between the query and the database.

15 16. The computer-implemented method of claim 1, wherein production of the unified search result shows speedup.

17. The computer-implemented method of claim 1, wherein the collecting of the search results and the generating of the unified search result are performed by a single computer.

18. The computer-implemented method of claim 1, wherein the collecting of the search results and the generating of the unified search result are performed by multiple members
20 of the group of computers operating in parallel.

19. The computer-implemented method of claim 1, wherein the collecting of the

search results and the generating of the unified search result of step (d) is performed by interleaving the search results from the executed searching tasks on the basis of numerical raw scores generated during the executed searching tasks.

20. The computer-implemented method of claim 1, wherein the total search to be performed is represented by a rectangular array, wherein each search task assigned to one of the group of computers operating in parallel is represented by adjacent rectangles making up the array, and wherein the total size of the array corresponds to the total search to be performed.

21. The computer-implemented method of claim 20, wherein each column of the array represents one of the sub-databases and each row of the array represents one or more queries.

22. The computer-implemented method of claim 21, wherein a width of at least one column of the array corresponds to a group of one or more sub-databases that can simultaneously be held in memory in one of the group of computers operating in parallel performing one of the searching tasks associated with that column of the array.

23. An apparatus for performing a computer-implemented method of searching a plurality of queries against at least one database containing a plurality of records, comprising:

- a. means for partitioning the plurality of queries into a set of smaller subsets of queries;
- b. means for partitioning the at least one database into a set of smaller subdatabases;
- c. means for designating searching tasks to be performed by

associating each of said subsets of queries with one or more of said subdatabases;

d. means for dividing at least one of the searching tasks into two or more smaller searching tasks, and designating the two or more smaller tasks as related tasks on a virtual shared memory bulletin board;

5 e. means for assigning each searching task to one of a group of computers operating in parallel, wherein each member of the group of computers operating in parallel has at least one searching task assigned thereto;

f. means for executing at least some of the assigned searching tasks using the group of computers operating in parallel;

10 g. means for collecting search results from the executed searching tasks; and

h. means for generating a unified search result in accordance with the collected search results;

wherein the means for partitioning of the queries and the means for partitioning of
15 the database comprise one or more members of the group of computers operating in parallel.

24. The apparatus of claim 23, wherein each member of the group of computers operating in parallel is identical.

25. The apparatus of claim 23, wherein at least two members of the group of computers operating in parallel are different.

20 26. The apparatus of claim 23, wherein each member of the group of computers operating in parallel has the same operating system.

27. The apparatus of claim 23, wherein at least two members of the group of computers operating in parallel have different operating systems.

28. The apparatus of claim 23, further comprising means for compensating for failure of one or more of the members of the group of computers operating in parallel, wherein said failure does not affect the correctness of the search results.

29. The apparatus of claim 23, wherein the means for executing at least some of the assigned search tasks further comprises means for reporting a numerical raw score as part of the search result associated with such task, wherein the numerical raw score corresponds to a quantitative measure of a match between the query and the database.

30. The apparatus of claim 23, wherein the means for collecting of the search results and the means for generating of the unified search result of step comprises means for interleaving the search results from the executed searching tasks on the basis of numerical raw scores generated during the executed searching tasks.

31. The computer-implemented method of claim 1, wherein step (c) further comprises dividing at least one of the searching tasks into two or more smaller searching tasks, and designating the two or more smaller tasks as related tasks on a virtual shared memory bulletin board.

32. The apparatus of claim 23, further comprising means for dividing at least one of the searching tasks into two or more smaller searching tasks, and designating the two or more smaller tasks as related tasks on a virtual shared memory bulletin board.